

Verification of railway interlocking - Compositional approach with OCRA

Christophe Limbrée¹, Quentin Cappart¹, Charles Pecheur¹, and
Stefano Tonetta²

¹ Université catholique de Louvain, Louvain-La-Neuve, Belgium
{quentin.cappart|christophe.limbree|charles.pecheur}@uclouvain.be

² Fondazione Bruno Kessler, Trento, Italy
tonettas@fbk.eu

Abstract. In the railway domain, an electronic interlocking is a computerised system that controls the railway signalling components (e.g. switches or signals) in order to allow a safe operation of the train traffic. Interlockings are controlled by a software logic that relies on a generic software and a set of application data particular to the station under control. The verification of the application data is time consuming and error prone as it is mostly performed by human testers.

In the first stage of our research [?], we built a model of a small Belgian railway station and we performed the verification of the application data with the NUSMV model checker. However, the verification of larger stations fails due to the state space explosion problem. The intuition is that large stations can be split into smaller components that can be verified separately. This concept is known as compositional verification. This article explains how we used the OCRA tool in order to model a medium size station and how we verified safety properties by mean of contracts. We also took advantage of new algorithms (k-liveness and ic3) recently implemented in NUXMV in order to verify LTL properties on our model.

1 Introduction

In the railway domain, an interlocking is a signalling subsystem that controls the routes, the switches and the signals before allowing a train through a station. Computer-based interlockings are configured based on a set of application data particular to each station. The safety of the train traffic relies on the correctness of the application data. Usually, the application data are prepared manually and are thus subject to human errors. For example, some prerequisites to the clearance (e.g. green light) of the origin signal of a route can be missing. This kind of error can easily be discovered by a code review or by testing on a simulator. However, errors caused by concurrent actions (e.g. route commands) are much harder to find. In this case, the combination of possible concurrent actions explodes quickly and testing all possible combinations manually is impracticable. The goal of our research is to develop a method based on model checking in order to verify the application data. Especially, our approach must scale-up and allow the verification of real size interlocking areas.

In a previous work [?], we built a model of a small Belgian railway station and we performed the verification of the application data with the NUSMV model checker. However the verification of larger stations fails due to the state space explosion problem: the models are too big so that the model checker does not give a result in reasonable time. In this paper, we therefore tackle the problem with a compositional approach. The intuition is that large stations can be split into smaller components that can be verified separately. We report on the usage of the OCRA tool in order to model a medium size station and on how we verified safety properties by mean of contracts. We also took advantage of new algorithms (k-liveness and ic3) recently implemented in NUXMV in order to verify LTL properties on our model.

Outline The paper is structured as follows. In *Section 2*, we give a brief overview of the formal techniques that have been used in the case study. In the *Section 3*, we describe our model and the new features that we have added compared to our first model. In *Section 4*, we explain our verification strategy for larger stations. In *Section 5*, we discuss the performance of our verification approach and show how counter examples are produced when we insert errors in the application data. References to related work are provided in *Section 7*.

2 Contract Based Verification

2.1 Symbolic Model Checking

Model checking [?] is a method to formally verify that a system is correct. In symbolic model checking [?], a system M is described by a finite set V of variables, the initial states are represented by a formula I over V , while the transitions by a formula T over the variables V and V' , where V' represent the value of V after a transition. In the scope of this paper, we consider finite-state systems. Thus, without loss of generality, we can consider V as Boolean variables and formulas in propositional logic.

A state is an assignment to the variables in V . An initial state is a state that satisfies I . A transition is a pair of states that satisfy T . A path is a sequence $\sigma = s_0, s_1, s_2, \dots$ of states such that s_0 is an initial state ($s_0 \models I$) and, for every $i \geq 0$, s_i, s_{i+1} is a transition ($s_i, s_{i+1} \models T$). A state s is reachable if there is a path s_0, s_1, s_2, \dots such that $s = s_i$ for some $i \geq 0$.

In this paper, we specify transition systems in SMV [?], the input language of different model checkers such as NUSMV [?] and NUXMV [?]. Safety properties have been formalized by invariants, i.e. formulas over V that must be satisfied by all reachable states. Temporal properties have been formalized into LTL [?], which uses temporal operators to specify the temporal evolution of the transition system. The typical LTL formula we consider is in the form $\mathbf{G} (\phi_1 \rightarrow F\phi_2)$, where ϕ_1 and ϕ_2 are state formulas over V . It means that whenever ϕ_1 is true along an execution, ϕ_2 is true in a state that follows along the trace.

2.2 nuXmv: Verification of Components with K-Liveness and IC3

In the case study we use NUXMV to prove invariants and LTL properties. In particular, we use the IC3 algorithm to prove invariants and the k-liveness algorithm for LTL properties.

IC3 [?] is a SAT-based algorithm for the verification of invariant properties of transition systems. Very briefly, the idea of IC3 is to build iteratively a sequence of formulas F_0, F_1, \dots, F_k such that i) $F_0 = I$, ii) for all $i > 0$, F_i is a set of clauses, iii) $F_i \models F_{i+1}$, iv) $F_i(V) \wedge T(V, V') \models F_{i+1}(V')$, and v) for all $i < k$, $F_i \models P$ where P is the property that we want to verify. The formulas F_i are therefore over-approximations of the state space reachable in up to i transitions. They are iteratively strengthened and extended by generalizing clauses while disproving candidate counterexamples. The procedure terminates when either a counterexample is found or when $F_i = F_{i+1}$ for some i so that F_i is an inductive invariant that proves P .

In [?], IC3 has been integrated with *predicate abstraction* (PA) [?]. The approach leverages *Implicit Abstraction* (IA) [?], which allows to express abstract transitions without computing explicitly the abstract system, and is fully incremental with respect to the addition of new predicates.

k-liveness [?] reduces liveness to a sequence of invariant checking. It uses a standard approach to reduce LTL verification for proving that a certain signal f is eventually never visited ($\mathbf{F G} \neg f$). The key insight of k-liveness is that, for finite-state systems, this is equivalent to find a K such that f is visited at most K times, which in turn can be reduced to invariant checking. k-liveness is therefore a simple loop that increases K at every iteration and calls a subroutine safe to check the invariant. In particular, the implementation in [?] uses IC3 as safe and exploits the incrementality of IC3 to solve the sequence of invariant problems in an efficient way.

2.3 OCRA: Contract-Based Compositional Approach

In this paper, we adopt a compositional contract-based approach and we use the framework supported by the OCRA tool [?]. In particular, we specify component interfaces in terms of Boolean data ports and LTL contracts.

The OCRA input language is a component-based description of the system architecture where every component is associated with one or more contracts. Each contract consists of an assumption and a guarantee specified as LTL formulas. The assumption represents a requirement on the environment of the component. The guarantee represents a requirements for the component implementation to be satisfied when the assumption holds.

When a component S is decomposed into subcomponents, the contract refinement ensures that the guarantee of S is not weakened by the contracts of the subcomponents while its assumption is not strengthened. This is checked independently from the actual implementation of the components and is verified by means of a set of proof obligations in LTL, which are discharged with model checking techniques [?].

OCRA allows to associate to a component a behavioral model representing its implementation. The language used for the behavioral model is SMV. OCRA checks if the SMV model is a correct implementation of the specified component simply calling NuSMV to verify if the SMV model satisfies the implication $A \rightarrow G$ for every contract $\langle A, G \rangle$ of the component.

3 System and Model Description

In this section, we describe the station, the model, and two new features of our model that are the directional locking and the sequential release.

3.1 The Station

Braine l'Alleud station, shown in Fig. 1, is a medium size Belgian railway station comprising 32 routes, 12 switches, 12 signals, and 4 platforms (101-104). A platform is a section of pathway, alongside rail tracks at a railway station, metro station or tram stop, at which passengers may board. A route is a line of railway track between two signals on a rail system (e.g. route $R_{CC.102}$ from signal CC to track 102 - signal JC). The station can be decomposed into two separate nearly symmetrical parts comprising 16 routes each: $M1$ and $M2$.




Fig. 1: Track layout of Braine station

3.2 Composite System

The two parts of the station (i.e. $M1$ and $M2$) are not totally independent but have interfaces. These interfaces materialize a mutual exclusion mechanism preventing two trains to head for a platform in opposite direction at the same time (e.g. routes $CC.101$ and $KC.101$). Such routes are called conflicting routes. The exercise then consists in defining the system and its components, the interaction among the components, and try to prove some global properties on the system by making assumptions on the environment of each component.

The cuts (i.e. $M1$ and $M2$) are chosen so-that: 1) the number of interface variables is minimum, and 2) it sticks to the principle of distribution between

interlockings applied in larger stations. The same principle will be applied to two interlockings sharing a section. As shown in Fig. 2, the system is made of three components: $M1$, $M2$, and $C1$. Partial Listing 1.1 shows how the components, and the interfaces are defined in OCRA. The components: $M1$, $M2$, and $C1$ are implemented in SMV language.

```

1 COMPONENT BraineLL system
2
3 REFINEMENT
4
5     SUB BraineLeft   : M1;
6     SUB BraineRight  : M2;
7     SUB Controller   : C1;
8
9 CONNECTION BraineLeft.BSIB_101 := BraineRight.BSIB_101;
10 ...
11
12 COMPONENT M1
13
14 INTERFACE — From Environment
15     INPUT PORT BSIB_101: boolean;
16     ...
17
18 COMPONENT M2
19
20 INTERFACE — From Environment
21     OUTPUT PORT BSIB_101: boolean;
22     ...

```

Listing 1.1: System definition in OCRA

The components are defined by means of the *SUB* keyword. The interfaces are defined as *INPUT* or *OUTPUT* (e.g. BSIB_101 is an output for $M1$ and an input for $M2$). The *INPUT* and *OUTPUT* are connected by mean of the *CONNECTION* keyword.

Figure 2 shows how the components are connected by interfaces. The L_CS OUTPUT variable (=TRUE) is an output of the system and states that the route is set and the origin signal at proceed aspect (e.g. the route R_{CC_102} is set and signal CC is green). The Controller outputs the cmdR variable stating that the controller has issued a route command. The Rongo{1,2} INPUT variable provides an acknowledgement that the route command has been properly processed by the interlocking. The two $M1$ and $M2$ interlocking components exchange the state of the platform track-circuits and the state of the BSI. A track-circuit is an electrical circuit that detects the presence of train in a block of track. The four track-circuits at the platform can be occupied by a train running in either $M1$, or $M2$. The BSI variable allows for mutual exclusion of conflicting routes leading to the same platform. The principle of functioning of the *BSI* is explained in *Sect. 3.4*.

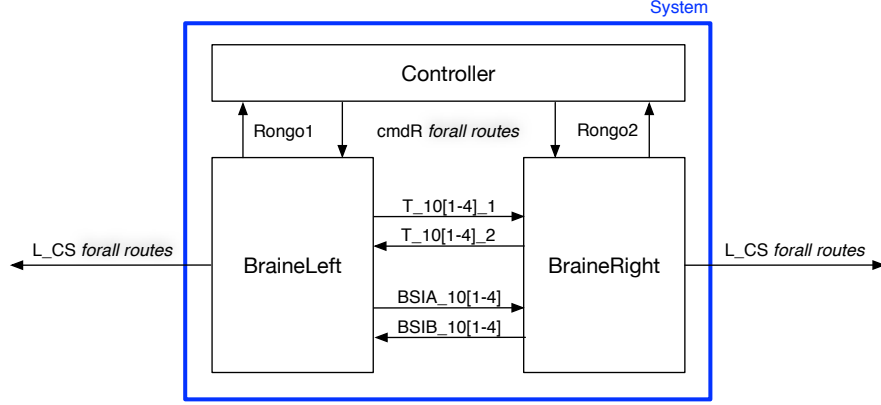


Fig. 2: Architecture of the composite system

3.3 M1 and M2 models

Figure 3 depicts the internal architecture of the $M\{1,2\}$ component. Each component is implemented in an SMV model. All the modules represent a function achieved by the interlocking except for the train module. In fact the train module allows to simulate the interact of the interlocking with its environment.




Fig. 3: Architecture of the SMV interlocking model

The interlocking module is directly translated from the application data by mean of a translator tool described in [?] and models the routes and the locking logic of the switches. Upon a route request, the interlocking (1) verifies that the route can be set and then controls the track components accordingly. A proceed aspect (e.g. green) is sent to the origin signal of the route when the switches are locked in correct position and the track-circuits are clear (i.e. no

other train is present on the route). Finally the interlocking detects the trains movement, releases the route and unlocks the resources used by the route. The track components (2) record the status of the track-side objects. For example: for a switch upon a command, the instance verifies that it is not locked before allowing the transition from one position to the other (e.g. left to right). The train modules (3) rely on the track layout of the station. When a signal is at proceed aspect, it simulates a train movement by actuation of the track components. This module is built independently of the application data by mean of a DSL (Domain Specific Language). The train module is local to $M\{1,2\}$ as it is built based on the track layout of its component.

3.4 BSI Interface Explained

In order to prevent head to head train collisions, the interlocking use a locking mechanism (i.e. *BSI* - Blocage du sens intermittent in French) that prevent two train to head for the same platform in opposite direction. Fig. 4 illustrates how the *BSI* variables are actuated upon a route command.

For each platform, two locking variables are used (e.g. *BSIA_102* and *BSIB_102* for platform 2). When no route is set towards platform 102 , the two variables have a permissive value (Free). Upon a route command (e.g. *R_CC_102*), the *BSIA_102* variable is set in a restrictive state (Locked). The routes in opposite direction (e.g. *KC_102*) are thereby blocked and the signal *KC* can never be commanded to a proceed aspect (e.g. green). The *BSIA_102* variable regains its permissive value when the train has reached platform 102.




Fig. 4: Directional locking for platform 102

3.5 Sequential Release

When the interlocking grants access to a route, it locks all the resources that will be run through by the train: typically all the switches and the track-circuits.

This prevents different routes that share those resources to be set at the same time. Such routes are called conflicting routes. Normally those resources are unlocked when the train has completely run through the route. They then become available for other routes. In large stations, it might be interesting to unlock the resources sequentially allowing them to be used by other routes before the train has totally run through. This contributes to improve the train traffic.

The principle of sequential release is illustrated in Fig. 5: the first route R_DXC_091 is set and prevents the second route R_DXC_092 to be set. The following switches are locked: P1A: left, P2B: right, and P3: right. According to the sequential release principle, the second route is set when the train T_2 is on the track-circuit T_01AC and when the track-circuit T_02BC is free.




Fig. 5: Sequential release example

4 Verification

The decomposition of one interlocking into several components allows to perform the verification on smaller models (one for each component) and thus limits the so-called state space explosion problem. Therefore we have used two different methods to verify the application data of Braine station: the first takes advantage of the OCRA compositional verification tool and the second uses the NUXMV tool to verify local properties.

The compositional verification applies to the safety properties that imply an interaction between the two components. Those properties are expressed by mean of contracts. An example of contract is given in Listing 1.2.

A second set of properties are verified straight on each component (i.e. $M1$ and $M2$) with NUXMV. Several instances of NUXMV can be started at the same time in order to reduce the computation time of the verification.

4.1 Compositional Verification

Conflicting Routes Controlled by Two Different Components The routes R_CC_101 and R_KC_101 are conflicting because they share the same platform as a destination and the corresponding safety property is expressed by the formula: $P = G!(R_CC_101_LCS \ \& \ R_KC_101_LCS) - routesTowards_101$ in Listing 1.2. Equation (4.1) shows how this property is verified by composition of the $M1$ and $M2$ modules. The first premise states that when the route R_CC_101 is set and origin signal is clear (i.e. $R_CC_101_LCS$ is true), the mutual exclusion property ($!BSIA_101 \ \& \ BSIB_101$) is true. The second premise states the same property (P_2) for the route R_KC_101 . P_1 and P_2 are respectively CC_101_OK and KC_101_OK in Listing 1.2. OCRA performs the verification of these two properties with NUXMV. The third premise states that the first two premises entail the global property P . Finally when all three premises are true, the composition of the components $M1$ and $M2$ verifies the global property P .

In other words, if each component (i.e. $M1$ and $M2$) properly blocks the access to a shared platform when it controls a route, then the other component will not be able to control a conflicting route for the same platform.

$$\frac{\begin{array}{l} \text{(Premise 1) } M_1 \models P_1 \\ \text{(Premise 2) } M_2 \models P_2 \\ \text{(Premise 3) } P_1 \wedge P_2 \models P \end{array}}{M_1 \parallel M_2 \models P}$$

Equation 4.1: Compositional verification of conflicting routes property involving the M_1 and M_2 components

```

1  CONTRACT routesTowards_101
2    assume: always TRUE;
3    guarantee: always (R_KC_101_LCS -> !R_CC_101_LCS );
4
5  CONTRACT routesTowards_101
6    REFINEDBY M1.CC_101_OK, M2.KC_101_OK;
7
8  CONTRACT CC_101_OK
9    assume: TRUE;
10   guarantee: always (R_CC_101_LCS -> (!BSIA_101 & BSIB_101));
11  CONTRACT KC_101_OK
12    assume: TRUE;
13   guarantee: always (R_KC_101_LCS -> (!BSIB_101 & BSIA_101));

```

Listing 1.2: Contract definition for conflicting routes towards platform 101 involving the *M1* and *M2* components

Listing 1.2 illustrates how the conflicting routes contract for the routes *R_KC_101* and *R_CC_101* is specified. First a top level contract (*routesTowards_101*) specifies that the two routes cannot be set at the same time. The top level contract is then refined by two contracts that apply on *M1* and *M2*: *KC_101_OK* and *CC_101_OK* respectively. These two contracts allow to verify that the *M1* and *M2* components handle the *BSI* locking mechanism properly. The syntax of the language is given in ([?]).

4.2 Local Safety Properties

The term *Local Properties* designates the properties that are not influenced by the environment of the component on which they are verified. Those properties are verified on each component SMV model with NUXMV. Due to the space limitation, those properties will not be explained in detail but examples are provided in Listing 1.3. They are expressed in two different ways:

- By mean of invariants (lines 1 to 5)
- By mean of *LTL* formulas and especially by using the ic3 algorithm (lines 6 and 7)

```

1 check_invar -p "!(M1.t1.front = derailed)"
2 check_invar -p "!(M1.t1.front = M1.t2.front)"
3 check_invar -p "!(M1.T_01AC.st = o) & M1.P_01AC.willMove)"
4 check_invar -p "(M1.R_CXC_103.L_CS -> !M1.R_EC_091.L_CS)"
5 check_invar -p "(M1.f1.U_CXC_13C.st = 1 -> (M1.f1.U_13C_15C.st
    = 1 xor M1.f1.U_13C_DXC.st = 1))"
6 check_ltlspec_klive -p "G (M1.U_IR_01AC.st = 1 -> ((M1.P_01AC.
    posi = cdr -> X M1.P_01AC.posi = cdr) & (M1.P_01AC.posi =
    cdn -> X M1.P_01AC.posi = cdn)))"
7 check_ltlspec_klive -p "G((M1.T_01AC.st = o & M1.TRP_CC.krc =
    s) -> X (!M1.R_CC_101.L_CS & !M1.R_CC_102.L_CS & !M1.
    R_CC_103.L_CS & !M1.R_CC_104.L_CS))"

```

Listing 1.3: Local properties

Explanation of the properties:

- Line 1: the train never derails. A derailment happens when a train takes a trailing point in reverse direction.
- Line 2: two trains never collide. This is done by verifying that their front never reaches the same track segment at the same time.
- Line 3: a point never move when its home track-circuit is occupied.
- Line 4: conflicting routes are not set at the same time. This formula verifies the same property as the contracts defined in OCRA.

- Line 5: the sub-routes are released in the correct sequence.
- Line 6: a point never moves when its latching variable is in restrictive state. These formulas are checked by mean of k-liveness (see [?])
- Line 7: signal replacement. The origin signal of a route is immediately commanded to red (replaced) when the train occupies the first track-circuit of the route and has triggered the first passage sensor. This prevents a second train to use the same authorization (i.e. signal green).

5 Results and Performance

In this section, we discuss the performance of our verification approach based on composition and local verification. We also illustrate how we validate the model and the properties by error seeding.

5.1 Performance

The tests were performed on 2.3 GHz i7 MacBookPro with 4 GB of RAM running under OS 10.11. Tables 1, 2, and 3 illustrate the results (in terms of computation time), which we achieve using different methods and different models. “BDD” refers to the fix-point algorithm using BDDs (see [?]); “SAT(ic3)” refers to the ic3 algorithm using a SAT solver as backend (see [?]); “SMT(ic3)” refers to the ic3 algorithm integrated predicate abstraction using an SMT solver as back-end (see [?]).

Table 1: Performance of the verification of invariants on monolithic models

Model	Tool	Method	Properties	Duration
1 Monolithic model	NuSMV	BDD	Invariants	$> 1 \text{ day}$
2 Partial monolithic model	NuSMV	BDD	Invariants	$> 1 \text{ day}$
3 Monolithic model	NuXMV	SAT(ic3)	$10 \times$ Invariants	123 sec
4 Monolithic model	NuXMV	SMT(ic3)	$10 \times$ Invariants	80 sec

Table 1 reports the performance of the verification of the application data for the station described in Section 3. Line 1 shows that NUSMV could not terminate in one day. After reducing the size the state space by allowing only 16 routes to be commanded, NUSMV could build the reachable state space in 6 days and verify invariants (line 2). One of the features of OCRA is to allow to rebuild a monolithic (32 routes) model based on the definition of the system. The verification of invariants is therefore possible. The results clearly show that ic3 with predicate abstraction performs better than plain ic3, and that both outperform the BDD-based algorithm on this case study.

Table 2: Performance of the verification of the contracts by OCRA

Model	Tool	Method	Properties	Duration
5	Contract refinement	OCRA	-	$4 \times$ Contracts <i>7,34 sec</i>
6	Implementation M1	OCRA	ic3	$4 \times$ Contracts <i>5,6 sec</i>
7	Implementation M2	OCRA	ic3	$4 \times$ Contracts <i>14,94 sec</i>
8	Composite monolithic	OCRA	ic3	$4 \times$ Contracts <i>1242 sec</i>

Table 2 illustrates the performance of the verification of the contracts and their implementation. Line 5 corresponds to the verification of the premise 3 of Equation (4.1) ($P_1 \wedge P_2 \models P$). Lines 6 and 7 are respectively related to the verification of the premisses 1 and 2 ($M_1 \models P_1$ and $M_2 \models P_2$). The sum of the duration of these 3 tasks gives the time needed by OCRA to check the coherence between the contracts and their implementation in the SMV models (i.e. ≤ 28 sec). This time is to be compared with the *1242 sec* needed by OCRA to verify the same contracts and implementations on a monolithic model.

Table 3: Performance of the verification of the local properties

Model	Tool	Method	Properties	Duration
9	M1	NuXMV	BDD	$197 \times$ Invariants <i>123 sec</i>
10	M2	NuXMV	BDD	$199 \times$ Invariants <i>424 sec</i>
11	M1	NuXMV	SAT(ic3)	$12 \times$ LTL <i>960 sec</i>
12	M1	NuXMV	SMT(ic3)	$12 \times$ LTL <i>20 sec</i>
13	M2	NuXMV	SAT(ic3)	$12 \times$ LTL <i>1036 sec</i>
14	M2	NuXMV	SMT(ic3)	$12 \times$ LTL <i>740 sec</i>

Finally Table 3 illustrates the verification of the local safety properties on the M1 and M2 components. Two approaches are used: first some invariants are verified with NUXMV and the standard BDD and second 12 LTL properties are verified with the ic3 algorithm. An order file based on [?] is used to optimize the BDD structure. ic3 with abstraction and the SMT MathSAT solver outperforms ic3 with MiniSAT in this context in an order of magnitude close to 50.

5.2 Error Seeding

In order to gain confidence in our model and properties, we have seeded errors in the model by removing some safety conditions in the route proving conditions³. As expected, OCRA could not prove the safety property and produced a counterexample. Listing 1.4 shows that the property is false (line 1) because the route R_KXC_101 is set (line 30) whereas the BSIB_101 is TRUE (line 5).

³ Conditions to give a proceed aspect on origin signal of the route.

<pre> 1 LTL spec G (R.KXC_101_LCS -> 17 (!BSIB_101 & BSIA_101)) id8 false 2 Trace Description: IC3 counterexample 3 -> State: 2.1 <- 4 ... 5 BSIB_101.st = TRUE 6 ... 7 -> Input: 2.2 <- 8 cmdR = R.KXC_101 9 -> State: 2.2 <- 10 R.KXC_101.cmd = TRUE 11 ... 12 -> Input: 2.3 <- 13 cmdR = R.KXC_103 14 -> State: 2.3 <- 15 R.KXC_101.st = s 16 ... </pre>	<pre> -> State: 2.4 <- U_IR_09C.st = 1 U_IR_07BC.st = 1 BSIA_101 = TRUE U_IR_07AC.st = 1 U_16C_JXC.st = 1 U_18C_16C.st = 1 U_KXC_18C.st = 1 25 ... 26 -> State: 2.5 <- R_KXC_101.st = rsu T_101.st = c T_101_1 = FALSE R_KXC_101_LCS = TRUE (Route is set) 31 KXCopen = TRUE </pre>
---	---

Listing 1.4: Error trace generated after error seeding in the model

6 Related Work

Many works applied model checking to interlocking systems. One of the first work dates back to 1998 and is described in [?]. However, as also concluded in [?], although small scale interlocking systems can be addressed by model checking, interlockings that control medium or large railway needs to tackle the state-space explosion problem. As shown also in [?], a single approach is often not sufficient to prove all properties and sometimes a combination of approach may dramatically improve the performance.

Compositional approach is one method to reduce the complexity of the verification but is not the only one. For instance, Cappart et al. [?] introduced a method based on discrete event simulations. The idea is to do not verify all the states but to limit the verification to a set of likely scenarios. However this method does not provide enough confidence that all the errors in the application data will be detected.

In [?,?], Winter shows how to compute optimized variable and transition orderings in order to speed-up the symbolic model checking of railway interlockings with NuSMV. She also reported on her findings on how to set the threshold for cluster.

In [?], Winter et al. modelled the interlocking by means of the formal notation ASM that are more readable. The formal model is translated in NuSMV code and the Safety requirements are expressed in CTL.

In [?], Peter Duggan (Siemens Rail Automation, UK) and Arne Borälv (Prover Technology AB, Sweden) have demonstrated that the Prover⁴ tools were

⁴ <http://www.prover.com>

successfully used to generate and test the configuration data of a realistic size UK station.

In [?], Haxthausen et al. detailed how they modelled an ETCS level 2 compatible Danish interlocking with the RT-Tester. The state space, the transition relation and the safety properties are efficiently evaluated by the SMT solvers that support bit vector and integer arithmetic. The model also include the sequential release feature.

In [?], Xu et al. verifies hybrid safety properties of Automatic Collision Avoidance System (ACAS) in the European Train Control System (ETCS). They verify those properties using Compositional Verification rules based on weakly monotonic time extension.

In [?], Antoni et al. have developed a SIL4 interlocking that uses the Petri Nets as application data. In [?], Dutilleul et al. have also used the Petri Nets in order to define a model pattern of railway interlocking.

7 Conclusions and Future Work

Conclusions

The verification of medium and large interlocking data is still a challenge due to the state space explosion problem affecting the model checking process. Our main contribution was to achieve the verification of the application data of a medium size railway interlocking by mean of compositional verification. In order to do that, we modelled our case study interlocking as a composite of smaller interlocking components in OCRA and SMV language. The verification of the safety properties (expressed as contracts) was performed with OCRA and NUXMV tools.

We have also added the sequential release functionality into our interlocking model. This functionality allows to increase the throughput of the railway network by releasing the route components earlier.

Finally, we have achieved the verification of LTL properties in efficient time thanks to the usage of the new ic3 algorithm implemented into NUXMV. The verification of the local components can be paralleled by running several instances of NUXMV at the same time.

Future work

In our future work, we will continue to refine the structure of the interlocking composite into adequate components (e.g. train). Our goal is to be able to verify safety properties on a network of interlockings by mean of compositional verification.

We will continue to develop the automatic translator tool in order to convert the application data of a network of interlockings into OCRA language.

Another goal is to develop a model of an IL/ETCS installation in order to verify safety properties related to the train dynamic characteristics (i.e. speed and position). In order to do that we will extend our train module in order to make it continuous.